

Quarterly Report
January – March 2009

Building the Framework for the
National Virtual Observatory

NSF Cooperative Agreement
AST0122449



INTERNATIONAL VIRTUAL OBSERVATORY ALLIANCE



Executive Summary	1
Activities by WBS	2
1 Management	2
1.1 General (planning, reporting, communications, team meetings, etc.).....	2
1.2 Science	2
1.3 Technical (including standards, configuration management).....	2
1.4 Financial.....	2
1.5 International coordination/collaboration	3
2 Science Requirements	3
2.1 Usage scenarios for all areas of astronomy research, including theoretical simulations.....	3
2.2 Requirements analysis	3
2.3 Demonstration definition and review	3
3 Operations, System Integration, and Testing	3
3.1 Quality assurance and software engineering design.....	3
3.2 Facility operations	3
3.3 User support.....	4
4 Registries	4
4.1 Resource metadata.....	4
4.2 Resource metadata schema	4
4.3 Publishing and harvesting protocols.....	4
4.4 Search protocols.....	4
4.5 Replication, synchronization, maintenance, revision control, and curation	5
5 Data Models	5
5.1 High-level (image, spectrum, time series, event lists, visibilities, catalogs, simulations, data quality).....	5
5.2 Low-level (measurement, quantity, uncertainty, relationship)	5
5.3 Descriptors and ontologies (UCDs)	5
5.4 Space-Time and regions	5
5.5 Standard schema.....	5
6 Data Access Layer	5
6.1 Data access services (catalog, image, spectrum, time series, visibilities, ...)	5
6.2 Data representation (VOTable, etc.).....	7
6.3 Framework (mediators, components).....	7
6.4 Data provider/consumer implementations and end-to-end testing.....	7
7 Query Language	8
7.1 Low-level: Astronomical Data Query Language	8
7.2 Mid-level: VOQL and OpenSkyQuery/OpenSkyNode.....	8
7.3 High-level: Complex queries	8
8 Web and Grid Services	8
8.1 Web Services (SOAP, WSDL, etc.).....	8

8.2	Grid Services (OGSA).....	8
8.3	Computational resource management	8
8.4	Virtual data.....	8
8.5	Application and service integration with Grid.....	9
9	Applications.....	9
9.1	Data location services.....	9
9.2	Cross-correlation services.....	9
9.3	Image combination, registration.....	10
9.4	Visualization tools and services	10
9.5	Theory.....	10
9.6	Statistical analysis	10
9.7	Data mining, outlier identification	10
9.8	Interfaces to/from legacy software systems.....	10
10	Community Engagement.....	10
10.1	Documentation	10
10.2	Web site	10
10.3	Technical training initiatives.....	10
10.4	Advocacy	10
11	Education and Public Outreach.....	10
	Activities by Organization	11
	Caltech–Astronomy Department and Center for Advanced Computational Research (CACR).....	11
	Caltech–Infrared Processing and Analysis Center (IPAC)	11
	High Energy Astrophysics Science Archive Research Center (HEASARC).....	11
	Johns Hopkins University.....	12
	National Optical Astronomy Observatories (NOAO).....	12
	National Radio Astronomy Observatory (NRAO)	12
	Raytheon/ADC (University of Maryland and George Mason University)	13
	Smithsonian Astrophysical Observatory	13
	Space Telescope Science Institute	13
	University of Illinois-Urbana/Champaign/National Center for Supercomputer Applications (UIUC/NCSA).....	14
	Publications and Presentations.....	16
	Virtual Observatory Articles in the Popular and Technical Press.....	17
	Usage Logs.....	18
	Acronyms.....	21

**Building the Framework for the National Virtual Observatory
NSF Cooperative Agreement AST0122449
Annual Report**

Period covered by this report: 1 January – 31 March 2009

Submitted by: Dr. Robert Hanisch (STScI), Project Manager

Executive Summary

The NVO Project continued to operate on its remaining carry-forward funds from FY08, with the result that development activities have stopped in all but a few areas and support and operational work is continuing on a best-effort basis only. This will most likely be the last Quarterly Report. We will begin a formal close-out of the project and submit a final report through the standard channels.

The NVO data discovery portal was released on 19 February. This represents a major re-design of the NVO web page and provides greater connectivity between portal components (directory, inventory, etc.).

Progress on the Table Access Standard stalled for a period of about two months at the beginning of the quarter. The NVO team has been working with the IVOA Data Access Layer WG Chair and others in IVOA leadership to try to get this work back on track. As of this writing, e-mail discussions have resumed on the DAL e-mail list. In concert with work on the TAP protocol, efforts are also underway to prepare the VO Registry for including TAP services. Our goal is to reach consensus on the TAP standard and associated issues at the May Interop meeting in Strasbourg.

Other preparatory activities for the Strasbourg Interop meeting included preparation of agendas for the Applications WG (T. McGlynn) and Grid and Web Services WG (M. Graham).

Routine testing of VO services has been expanded and now includes regular testing of all cone search and SIAP services known from the NVO registry.

Activities by WBS

1 Management

1.1 General (planning, reporting, communications, team meetings, etc.)

Regular weekly telecons of the Technical Working Group (TWG) continued, led by J. Good (IPAC). Technical issues addressed include:

- Table Access Protocol. TAP issues have occupied the largest fraction of the discussions. General access to tabular data is at the heart of much of the advanced VO functionality. Progress over the years on defining an international standard for a DBMS-like query standard has been slow, though recently significant work has been done within the NVO to define and prototype a simple parameter-based subset of this functionality.
- NVO Portal. Final refinement and release of the Portal components has also occupied a significant fraction of the group's discussion. These tools are now operation and the group is monitoring use patterns.
- Monitoring. The operational monitoring being done at HEASARC of all VO services has become a staple discussion topic. Recently, the original simple service availability checks have been augmented with periodic service validity checks.

The Executive Committee also continues to meet weekly by telecon. No face-to-face team meetings were held, however, given the current funding situation.

Considerable testing of the EVO videoconferencing facility was done during the period with several weekly telecons held this way. S. Emery Bunn Wrote custom documentation for NVO team members to use coordinated with the EVO support/developers to learn more about the tool, wrote documentation to assist team members in using it, and provided direct support and assistance during TechWG telecons. Unfortunately, EVO has not worked terribly well for us.

1.2 Science

Nothing to report this quarter.

1.3 Technical (including standards, configuration management)

Progress on the Table Access Standard stalled again for a period of about two months. The NVO team has been working with the IVOA Data Access Layer WG Chair and others in IVOA leadership to try to get this work back on track. As of this writing, e-mail discussions have resumed on the DAL e-mail list, but the DAL leadership has rejected our suggestions for a series of telecons focused on resolving the key outstanding issues. Our goal is to reach consensus on the TAP standard at the May Interop meeting in Strasbourg.

1.4 Financial

Total project expenditures since inception now stand at \$13,857,250 compared to a budget of \$14,151,966, leaving a balance of \$294,716. We have \$73,444 in invoices received and pending payment, leaving a net of \$221,272. Total spending this quarter, including pending invoices, is \$462,295. Although it appears that significant funds remain,

most of this is held in two or three organizations and, we believe, is covering work already completed but not invoiced.

1.5 International coordination/collaboration

Senior members of the NVO project participated in the IVOA Executive telecon on 19 February engaged in discussions on other matters of common interest. T. McGlynn chairs the IVOA Applications WG and M. Graham chairs the IVOA Grid and Web Services WG, and both have planned technical sessions for the May Interop meeting in Strasbourg. R. Hanisch chairs the IVOA Data Curation and Preservation IG, but no sessions are currently planned for Strasbourg.

S. Emery-Bunn (Caltech) prepared an article for submission to the next IVOA Newsletter, briefly describing the new NVO web portal.

The IAU Virtual Observatories, Data Centers, and Networks Working Group will meet during the IAU General Assembly in Buenos Aires in August 2009.

2 Science Requirements

2.1 Usage scenarios for all areas of astronomy research, including theoretical simulations

No activities to report this quarter.

2.2 Requirements analysis

No activities to report this quarter.

2.3 Demonstration definition and review

No activities to report this quarter.

3 Operations, System Integration, and Testing

3.1 Quality assurance and software engineering design

No activities to report this quarter.

3.2 Facility operations

M. Preciado (HEASARC) continues to operate a suite of tests using the NAGIOS software system which checks every site providing VO services. Typically tests are run every hour. For sites which provide many services, only a representative subset of the services are checked to avoid undue use of the services. Currently about 100 tests are made on servers which host about 7,000 distinct queryable services. All portal elements are tested.

Preciado reviews the status of all sites on a daily and sends notifications to sites where requests are failing. Critical sites are informed immediately while less critical and non-NVO sites are normally queried if they have been down for several hours. During the quarter, 40 issues were referred to the responsible parties. Two of these remain open.

Preciado has begun the preliminary deployment of a much finer grained testing. This uses the validation tools developed by R. Plante (NCSA) and tests every registered cone search and SIA service. A record of all issues found in testing the service is maintained in a database and is queryable using a Web interface. A programmatic interface is planned. During the initial testing it appears that less than 30% of all services are fully compliant with the VO standards. However many of the issues found are relatively mi-

nor problems with how errors are handled or use of single precision floating point numbers where double precision is mandated.

Aside from occasional service downtimes there were no major interruptions to services at the NVO's major web sites at CACR, NOAO, CfA, HEASARC, IRSA, NCSA, STScI, or JHU.

VOEvent services continued nominal operations at the CACR (A. Drake, M. Graham, R. Williams).

Montage services continued nominal operations at IRSA.

Discussions have begun on how to finalize the NVO repository to ensure that all software developed for the NVO is properly delivered.

3.3 User support

All user queries (54 in this quarter) were reviewed by S. Emery Bunn (CACR) and directed to appropriate NVO team members. Responses and resolutions for all items were obtained.

S. Emery Bunn (CACR) edited and published the third NVO newsletter. The newsletter was released March 24 (see <http://www.us-vo.org/newsletter/003/>).

Currently there are 433 subscribers to the mailing list.

4 Registries

With the new NVO portal deployed and the project operating on remaining no-cost extension funding, our current efforts in the area of Registries are concentrated primarily on maintenance of current services and participation in IVOA standards efforts.

4.1 Resource metadata

Describing TAP services as well as other services that return data in tabular form in the registry is a primary area of focus today. Consequently, the roadmap for the VO-DataService metadata standard, the VOResource XML extension schema used to describe data and services in the registry, is tied closely with that of TAP. An important discussion that has been going on this quarter is communicating the types associated with the table columns. While we use VOTable to transmit results from TAP services, the data types supported by VOTable may not correspond well with the types used by the underlying database—particularly from the perspective of forming a query against the database. We now have a proposal for a core set of database-oriented types that would be supportable by TAP services, map-able into VOTable types, and enumerated in the service description interface. We plan to solidify this proposal by the Spring IVOA meeting.

4.2 Resource metadata schema

No activity during this quarter.

4.3 Publishing and harvesting protocols

No activity during this quarter.

4.4 Search protocols

No activity during this quarter.

4.5 Replication, synchronization, maintenance, revision control, and curation

T. Dower (STScI) has put some effort into reviewing problems with resource descriptions and communicating with their publishers. In anticipation of a more extensive operations phase, she has been establishing some re-usable practices for this curation. In particular, she created some documentation that can be incorporated into emails to publishers regarding common problems with resource descriptions.

5 Data Models

5.1 High-level (image, spectrum, time series, event lists, visibilities, catalogs, simulations, data quality)

J. Cant continued work on updating the Java library that implements the Spectrum Data Model. A test version of the library was sent to the ESAC-Madrid team for review. FITS format support still needs to be completed. J. McDowell coded some simple test applications to exercise the library.

5.2 Low-level (measurement, quantity, uncertainty, relationship)

No activity during this quarter.

5.3 Descriptors and ontologies (UCDs)

NVO team members M. Graham and K. Borne were on the Scientific Organizing Committee of the conference “Practical Semantic Astronomy,” which was held in Glasgow, Scotland, 2-5 March 2009.

5.4 Space-Time and regions

J. Cant continued implementation of the STC Java library. He has completed an STC-X to STC-S converter program and is working on the inverse. After that, we plan to develop a GUI to let an astronomer interactively create an STC descriptor.

A. Rots had discussions with M. Demleitner (GAVO) and D. Berry (AstroGrid) on STC-S and is working on an update to the Note.

5.5 Standard schema

While the SAO team continues work on the SED library, J. Mazarella and the IPAC team are waiting for this and further specification of the VO photometry/SED Data Model required for interfacing NED's internal SED data model to a VO standard SED data model; they have concentrated on other work in this quarter.

6 Data Access Layer

6.1 Data access services (catalog, image, spectrum, time series, visibilities, ...)

Catalog access. Progress on the Table Access Protocol (TAP) has been slow this quarter, with NVO development mostly on hold while we negotiate with our international partners on the IVOA TAP specification, development of which is currently being led by P. Dowler (CADC). The basic service interface established last fall is sound and has thus far remained largely unchanged. Some additional progress has been made on specific technical issues including how to specify database and table metadata (the TAP_SCHEMA), and case sensitivity in ADQL queries. Other key issues such as ADQL usage for advanced queries and the role of a simplified, reliable parameter-based query method as an alternative to the more complex ADQL-based technology have not yet been resolved.

The parameter based query is needed by NVO to support the most common science use cases such as spatial queries and typical filter-type access to astronomical catalogs, for use while development of the more complex ADQL-based technology goes forward. The parameter query will eventually provide a replacement for the very successful simple cone search interface, as well as provide optimized support for spatial queries including multi-position queries where the user uploads a table of object positions, all of which are processed in a single query. The more powerful but complex ADQL (SQL) query technology is also needed for NVO science use cases and continues to be a high priority in the longer term for development and prototyping, but it will likely take a year or more for this technology to mature.

The TAP specification will be a major topic of discussion at the IVOA Interop meeting in Strasbourg in May, where we hope to reach agreement with our international partners on some of these issues. Regardless of the extent of the IVOA agreements on TAP reached in May we expect to finalize a first working version of the TAP specification in time to support internal NVO development and applications for use while work continues in parallel on the IVOA TAP specification.

Spatial Footprints. Some discussion of how to use the parameter-based technology developed for spectral access (SSA) has taken place (G. Greene, A. Rots, D. Tody), with the conclusion that this is probably adequate for describing the outline of a simple region. ESO is using the NVO DALServer framework for spectral and image services and the issue of how to represent spatial footprints within such services came up again. As the issue is still unresolved they will continue to publish simple observation footprints with the technique they have used previously, using the DAL2 extension mechanism.

Spectral and time series data. Both the older prototype spectral services and the newer, IVOA standard SSA 1.0 services are now present in the second generation registry, which came on line earlier this year. Use of SSA services by client applications is starting to pick up now that such services can be registered and discovered automatically by client applications. A SSA 1.1 update is planned once real world testing with client applications is further along. The main feature planned for SSA 1.1 is addition of a “getCapabilities” operation to describe the capabilities provided by a specific SSA service instance.

Image data. A conceptual design for the Simple Image Access version 2 (SIAV2) interface was discussed last fall at the IVOA Interop in Baltimore (D. Tody, F. Bonnarel, J. Salgado). Good progress has been made since that time on the key issues which need to be detailed before a first working draft of the SIAV2 interface can be prepared. These key issues include work on how to represent multi-dimensional world coordinate systems with reference to both FITS WCS and STC (F. Bonnarel, D. Tody, A. Rots) and a data model for polarization (A. Richards), emphasizing primarily radio data, which constitutes the bulk of the multi-dimensional, sometimes polarimetric data in our early use cases for cube data access with SIAV2.

A. Richards and D. Tody met in Socorro in early April to discuss these issues as well as plans for SIAV2 prototyping, including on the fly imaging of interferometry data and access to spectral, polarimetric, and time cubes using SIAV2. Some preliminary discus-

sions have been held with ALMA, Arecibo, and Jodrell in connection with cube data access prototyping for SIAV2.

An effort has begun to improve the characterization of interferometric radio data (both visibility data and any derived images or cubes). In order to spatially cross correlate radio observations with those from OIR, X-ray, and so forth we are working to define a standard for specifying the circular field of view of a radio observation. An additional issue for interferometers is to specify the sampling of spatial frequencies in the observation or in any derived image. D. Tody, L. Sjouwerman, A. Richards, J. Crossley, and J. Benson have thus far contributed to this effort.

Complex data. The concept of the *generic dataset* and representing complex data as an aggregation of typed datasets in a generic dataset query is fundamental to the DAL architecture and family of data access interfaces. A related concept has been explored recently in connection with TAP. The proposed Parameter Query Language (PQL) would provide a generic means to pose a query for any type of dataset, similar to what the generic dataset concept already provides. The problem with the proposed PQL is that it provides no provision for access to specific types of virtual data (tables and multi-table tablesets, images, spectra, etc.) nor for description of complex data aggregates; PQL as proposed only provides for discovery of generic datasets via a common parameter-based query. Further discussion within the IVOA is needed to jointly explore this issue further and hopefully clarify the general model for data discovery and access in the VO context.

6.2 Data representation (VOTable, etc.)

A proposed VOTable 1.2 specification is currently under review. This is largely backwards compatible with earlier versions, including the VOTable 1.1 specification which is what is most widely used within VO currently. Some issues are still under discussion, including a change to the INFO element which is not backwards compatible and which could impact much existing software including many DAL services. The means used to describe the coordinate frame of table columns has also changed to be more STC compliant.

6.3 Framework (mediators, components)

The second (implementation) phase of the applications framework project began in January. A project whitepaper and conceptual design were delivered to OPTICON and NVO to close out the first phase of the project. A prototype parameter file implementation and associated draft specification were completed in February. Within the US, the applications framework project scope has expanded to include AURA and AUI as well as VAO. This group prepared a whitepaper on future astronomical software which was submitted to the Decadal Review panel in February.

6.4 Data provider/consumer implementations and end-to-end testing

NED, ESO, and NRAO currently have efforts underway to use the NVO DALServer service framework to implement data services, with the experience from these efforts helping to improve new versions of the DALServer framework.

Over the past three months, R. Ebert, O. Pevunova, and J. Jacobson have made substantial progress in implementation of a Simple Spectrum Access (SSA) interface into the NED spectra database using the DAL Server Toolkit. Jacobson has successfully updated the NED-specific Java code to accommodate the latest version of the DAL Server

Toolkit. Jacobson and Ebert have been in communication with D. Tody regarding some problems that have been encountered. Pevunova has completed the addition of observation time information to the spectrum metadata, where available; she is presently assigning UTYPEs and UCDs to the various metadata parameters, and completing the Description tags for each field. All of this work is being done on an internal development/test server available only within IPAC. A new server has been purchased to provide public access to the prototype SSAP service; this is scheduled to be online in June 2009.

M. Fitzpatrick implemented support for SSAP (Simple Spectral Access Protocol) services in the VOClient library and VO-CLI commandline tools. SSAP v1.0 services are fully supported, however the majority of SSA services are varyingly compliant with earlier versions of the protocol and are only partially supported at the moment.

7 Query Language

7.1 Low-level: Astronomical Data Query Language

Nothing to report this quarter.

7.2 Mid-level: VOQL and OpenSkyQuery/OpenSkyNode

The SuperComputing 2008 (SC08) HPC Storage Challenge, a competition showcasing applications that use approaches that effectively utilize the storage subsystem, was won by the GrayWulf team from JHU, Microsoft, University of Illinois at Chicago, the University of Hawaii, and Dell, Inc. led by Szalay (JHU), for their entry "[Storage Challenge GrayWulf: Scalable Clustered Architecture for Data-Intensive Computing](http://www.jhu.edu/~gazette/2008/08dec08/08superteam.html)." The winning entry (<http://www.jhu.edu/~gazette/2008/08dec08/08superteam.html>) featured a query that searched for quasars through 20 TB of SDSS Stripe 82 test data in 12 minutes, a task that would take several days on a monolithic non-partitioned dataset without a spatial index.

The JHU group has nearly finished loading ~25 TB of simulated Pan-STARRS Year 1 data into a partitioned GrayWulf system. Szalay, Budavari, Wilton, and Dobos have begun work on incorporating the GrayWulf technologies into Open SkyQuery. JHU has hired a summer intern who will develop a VOSpace/VOPipe implementation using CasJobs in the second quarter. This will also be integrated into Open SkyQuery in the future.

7.3 High-level: Complex queries

Nothing to report this quarter.

8 Web and Grid Services

8.1 Web Services (SOAP, WSDL, etc.)

Nothing to report this quarter.

8.2 Grid Services (OGSA)

Nothing to report this quarter.

8.3 Computational resource management

Nothing to report this quarter.

8.4 Virtual data

No activities to report this quarter.

8.5 Application and service integration with Grid

B. Baker (NCSA) has been exploring the use of OpenID, an open and widely used standard for managing single sign-on (SSO) to web portals. To that end, he completed a demonstration of an OpenID-based authentication service (<http://nvoauth-openid.ncsa.uiuc.edu/openid/>) that can replace our current Pubcookie-based service. Users would see very little difference between the current system and the OpenID-based system, though the latter would offer two main advantages. First, portals could support non-NVO OpenID names from other providers like Amazon. Second, OpenID, as an open standard, would make it possible to build an IVOA standard for interoperable portal-based SSO for all VO portal services.

9 Applications

9.1 Data location services

A key milestone for the NVO project was the release of the NVO data discovery portal and new home page on February 17, 2009. This directly involved team members from CACR, IPAC, JHU, STScI and the HEASARC.

Results from a review of portal services by astronomers external to the NVO were collated by S. Emery Bunn and acted on by the members of the Portal Team. Changes were primarily to ensure that users understood what was possible in given situations and better documentation of what the elements of the portal did. The Registry service was renamed the 'Directory' since that seemed to be less confusing. Significant testing of our services revealed some bugs which were rectified.

A sweep of portal services for consistency in style and usage helped make the system look more consistent. A set of icons was developed for the portal services.

The portal was integrated into the NVO home page, and the OpenSkyQuery service was added to the portal elements for this purpose although it is not entirely integrated with the other portal elements.

M. Fitzpatrick (NOAO) implemented support for SSAP (Simple Spectral Access Protocol) services in the VOClient library and VO-CLI command line tools. SSAP v1.0 services are fully supported, however the majority of SSA services are varyingly compliant with earlier versions of the protocol and are only partially supported at the moment.

Fitzpatrick also began implementation of a new C-based API that implements the client-side SAMP (Simple Applications Messaging Protocol) functionality. This API may be used both to send SAMP messages to other applications (e.g. to load an image or table) as well as to respond to messages sent by other applications. Like the VOClient library, this new SAMP library is meant to simplify much of the underlying technology for the user and will provide bindings for a variety of languages used in legacy software. The central message router in the SAMP architecture (the so-called 'Hub') will possibly be implemented at a later time, release of this new library is expected in the next quarter.

9.2 Cross-correlation services

Nothing to report this quarter.

9.3 Image combination, registration

Nothing to report this quarter.

9.4 Visualization tools and services

Nothing to report this quarter.

9.5 Theory

Nothing to report this quarter.

9.6 Statistical analysis

Nothing to report this quarter.

9.7 Data mining, outlier identification

Nothing to report this quarter.

9.8 Interfaces to/from legacy software systems

Nothing to report this quarter.

10 Community Engagement

10.1 Documentation

Issue No. 3 of the NVO Newsletter was distributed on 23 March. The several hundred subscribers to the NVO mailing list are notified when the Newsletter is released or other events of interest to the NVO user community occur. The Newsletter is available on-line at <http://www.us-vo.org/newsletter/003/>.

10.2 Web site

The website redesign described in the previous Quarterly Report was formally released on 19 February.

10.3 Technical training initiatives

No activities to report this quarter.

10.4 Advocacy

No activities to report this quarter.

11 Education and Public Outreach

EPO activities are dormant in the no-cost extension phase of work.

Activities by Organization

Caltech–Astronomy Department and Center for Advanced Computational Research (CACR)

S. Emery Bunn collated the comments of the science testers of the portal and configured the new portal website. She fielded and monitored responses for 54 messages sent to the NVO feedback address. She provided documentation and logistical support for using the EVO network-based videoconferencing tool. She published the third NVO Newsletter, posted news items to the NVO website, and is supporting preparation of the second IVOA Newsletter.

Caltech–Infrared Processing and Analysis Center (IPAC)

J. Good led the weekly NVO Technical Working Group telecons. Issues discussed include the TAP protocol, the NVO portal and monitoring services.

O. Pevunova continued to maintain NED services in the NVO Registry, including access logs provided daily. In response to changes in the Cone Search protocol and requests from users building new VO clients to NED server content, the NED services have been updated to accept the VERB (verbosity) parameter. R. Ebert, O. Pevunova, and J. Jacobson have made substantial progress in implementation of a Simple Spectrum Access (SSA) interface into the NED spectra database using the latest version of the DAL Server Toolkit (NRAO); we expect to deliver a public prototype service for testing by the NVO team by June 2009. Pevunova also participated in the weekly Technical Working Group telecons, and J. Mazzarella participated in the weekly Executive Committee telecons.

High Energy Astrophysics Science Archive Research Center (HEASARC)

HEASARC activities were primarily in two areas: portal development and validation of VO services. T. McGlynn led the portal development to its release in February and also was the principle developer of the DataScope and SimpleQuery tools. He also helped to extend the portal documentation.

M. Preciado continually monitored all VO services and informed appropriate parties of issues. He has also developed a framework for detailed checking of Cone and SIAP service for compliance with the VO standards.

T. McGlynn is Operations WBS lead for the NVO and a member of the NVO Executive committee. He also chairs the IVOA Applications Working Group. An RFC for the Simple Applications Messaging Protocol was initiated by McGlynn has been promoted to IVOA recommendation during the past quarter (though the vast majority of the work there was done by M. Taylor outside of the NVO).

T. McGlynn s involved in discussion of TAP, VOQL and other data access protocols.

HEASARC staff also participated in the NVO team meetings, IVOA meetings, and the weekly technical telecons as well as participating in the IVOA group discussions as appropriate.

Johns Hopkins University

A. Thakar continues to maintain the NVO web logs and monitor the harvesting at JHU on a daily basis, and to contribute quarterly and annual usage stats and charts for the NVO quarterly and annual reports. During this quarter, Thakar had to resume interrupted log harvesting due to disk space issues on the JHU Weblog server.

Thakar continues to provide triage and support for Open SkyQuery and other VO services at JHU. In this quarter, he updated the Open SkyQuery website to conform to the NVO application template. He also reprinted SkyNodes during a server failure at JHU.

Thakar prepared a job description for requisition and posting for a VO Data Scientist at JHU. Ani screened and selected a summer student intern to work on a VO related project at JHU during May through July 2009.

T. Budavari, along with R. Wilton and L. Dobos (Eotvos University) designed and prototyped a third-generation SkyQuery engine. Their approach addresses the issues of outliers, missing information, large number of dimensions and the vast amount of data by combining elements of robust statistics and recursive algorithms that provide improved eigensystem estimates step by step, utilizing all the information contained in the observations.

G. Fekete cleaned up the HTM Spherical Loops library, preparing using it for HTM Cover generation, as well as area and point inclusion computation. This library produces a more compact and well-formed ordered list of arcs in which consecutive arcs in the ordered list follow the natural order inherent in the outline, and internal arcs are eliminated.

J. Raddick continued to maintain the virtualobservatory.org public outreach website.

National Optical Astronomy Observatories (NOAO)

M. Fitzpatrick continued participation in the Portal Working Group improving the VO-CLI tools prior to a general release. Enhancements were made to the ranking algorithm for Directory search results.

Fitzpatrick implemented support for the Simple Spectral Access Protocol (SSAP) services in the VOClient library and the VO-CLI command line tools.

Fitzpatrick began development of a new C-based API that implements the client-side Simple Applications Messaging Protocol (SAMP) functionality.

J. Schneider started development of tests for the NVO web pages. He setup a Subversion repository for the testing scripts and has begun creating Fressia test suites to verify all web links (first level) in the NVO web pages. (<http://sourceforge.net/projects/fressia>)

National Radio Astronomy Observatory (NRAO)

NRAO participation in VO technology development continues, focusing on development of the new table access protocol (TAP) and the second generation image access facilities (SIAV2), both efforts being led by D. Tody within the US NVO. Both protocols will be discussed at the upcoming IVOA interoperability workshop in Strasbourg in May. TAP will enable advanced table analysis applications including large scale distributed cross matching and advanced table and database queries. SIAV2 will support access to spectral and time cubes from modern instruments and surveys, including ALMA and EVLA.

Within the NRAO VO program work has begun on facilities to enable correlation of VLA observations with those from outside observatories. A VLA observation summary catalog has been produced, and has been published to the VO as a cone search service, using the DALServer service framework under development within the US VAO. A client program based upon the US VAO VOClient interface is under development to enable correlation of VLA and external observations. We have also begun looking at integration of the VLA catalog into the VAO inventory service hosted at IPAC. The KMLNow application, used to overlay astronomical images within Google Sky, was presented at the AAS in January as well as in a SPIE Newsroom article (J. Crossley, R. Duplain, N. Radziwill).

The second (implementation) phase of the applications framework project began in January. A project whitepaper and conceptual design were delivered to OPTICON and NVO to close out the first phase of the project (P. Grosbol, D. Tody, et. al.). A prototype parameter file implementation and associated draft specification were completed in February (L. Paioro, D. Tody, et. al.). Within the US, the applications framework project scope has expanded to include AURA and AUI as well as VAO. This group prepared a whitepaper on future astronomical software which was submitted to the Decadal Review panel in February.

Raytheon/ADC (University of Maryland and George Mason University)

George Mason University (GMU) staff K. Borne was co-author on a talk presented at the January 2009 AAS meeting by one of the student prizewinners from the 2008 NVO Summer School; the talk was titled "Probing the Quasar Distribution within the Virtual Observatory." Borne was on the Organizing Committee for the March 2009 international conference "Practical Semantic Astronomy" in Glasgow Scotland, which was co-sponsored by the IVOA semantics group. Borne is on the Program Committee for a workshop in August 2009 at UMBC on Earth and Space Science Informatics. During this reporting period, Borne met with and began discussions with the MAST archive group at STScI to implement VO-enabling scientific data mining algorithms that Borne and the data-mining experts at UMBC have been developing over the past few years.

Smithsonian Astrophysical Observatory

G.Fabbiano continued serving in the Interagency Working Group on Digital Data (IWGDD) and the Smithsonian Institution Digitization Strategic Plan Committee (SI DigiSPC). She also participated in discussions about the VO at Caltech and IPAC, and is scheduled to be a participant in an upcoming CfA town hall on the VO.

A. Rots gave a presentation on VO infrastructure at the CfA VO presentation on Apr 14.

A. Rots attended the January AAS and will attend Hotwiring the Universe II in April in Santa Cruz.

In related work, Rots has circulated the first draft of FITS WCS Paper V, on Time.

Space Telescope Science Institute

T. Dower implemented additional search functions for the new registry to assist other developers in transitioning their old VO services to use the new registry. She adjusted much of the text in the registry and publishing web pages to reflect the Portal Group's decision to call the higher-level registry interfaces a 'Directory', and for other terminology

changes dependent on that one. She also streamlined the registry caching system written in December and otherwise cleaned up the registry code and help text based on feedback from Portal testing. Dower made several improvements to the registry publishing interface to allow simple form editing of resources with multiple interfaces and capabilities of different types. The publishing service has been moved to the same production server as all STScI-based registry services. She has also continued performing maintenance and management of the registry, including beginning automatic harvesting from the newly-updated ESAVO registry, providing support for users of registry services within the NVO, and assisting resource owners in updating their own resource records via the publishing service.

K. Gillies deployed Version 1.0 of the APT/SIA service on the STScI servers in January. N. Gaffney demonstrated a Google Sky plugin accessing the APT/SIA service. The feature set needed by APT is satisfied by Version 1.0. Work is now proceeding to expand the APT/SIA footprint service code base to take advantage of the prototype observation data model and to replace functionality required by the HLA. Versions 1.1 and 1.2 were delivered during February to address requests from the APT team. Version 2 is now nearing readiness. HLA access has been isolated to allow a shared C# library for use in multiple projects. The big change in version 2 is that footprints are produced based on spherical library database tables allowing access to all HLA data. Work should begin soon to add the prototype observation data model.

R. Hanisch prepared the NVO Quarterly Report for the period 1 Oct 2008 through 31 Dec 2008. The report is available in the NVO document repository, at <http://www.usvo.org/pubs/files/FY09Q11.pdf>.

Hanisch edited the IVOA Standards Process document to incorporate the recently agreed upon document version numbering scheme, and provided logistical support for the IVOA Executive telecon on 19 February.

University of Illinois-Urbana/Champaign/National Center for Supercomputer Applications (UIUC/NCSA)

R. Plante continues to chair the IVOA Registry Working Group. The main efforts in this area are the Registry Interfaces document and the VODataService metadata schema extension. He also has contributed to the discussions of the Table Access Protocol standard.

Plante continues to support standard service validation services, including the IVOA's Registry of Registries. The NVO Operations efforts this quarter have made extensive use of the validation services to comprehensively test all standard services registered in the VO and collating the results. As we continue to do this, we discover new compliance and testing issues, some requiring further updates to the validation services themselves.

Plante has been presenting lectures from the NVO Summers School to a graduate-level seminar in the University of Illinois Department of Astronomy. In addition to providing outreach to new potential users, this has provided additional testing of the new NVO Discovery Portal prior to its final release, with bugs being reported to the portal developers. Plante found the lectures presented worked well outside of the NVO Summer School setting; however, he did rely heavily on his experience with all four NVO summer schools in understanding how to present the material.

B. Baker continues to support single sign-on services for the NVO. Our most active supporting portals today are the NOAO NVO Science Portal (nvo.noao.edu) and the Dark Energy Survey Project Team Portal (des.ncsa.uiuc.edu). He has been prototyping an OpenID-compatible version of our login services.

Publications and Presentations

Budavári, T., Wild, V., Szalay, A. S., Dobos, L., Yip, C.-W. 2009. MNRAS, 394, 1496. “Reliable eigenspectra for new generation surveys.”

Budavári, T. 2009. ApJ, 695, 747. “A unified framework for photometric redshifts.”

Budavári, T., Heinis, S., Szalay, A. S., Nieto-Santisteban, M., Gupchup, J., Shiao, B., Smith, M., Chang, R., Kauffmann, G., Morrissey, P., Schiminovich, D., Milliard, B., Wyder, T. K., Martin, D. C., Barlow, T. A., Seibert, M., Forster, K., Bianchi, L., Donas, J., Friedman, P. G., Heckman, T. M., Lee, Y.-W., Madore, B. F., Neff, S. G., Rich, R. M., Welsh, B. Y. 2009. ApJ, 694, 1281. “GALEX-SDSS catalogs for statistical studies.”

Virtual Observatory Articles in the Popular and Technical Press


K. Becker, 13 February 2009, Boulder Daily Camera. “Old data, fresh eyes.”
<http://www.dailycamera.com/news/2009/feb/13/science-old-data-fresh-eyes/>

J. Crossley, R. DuPlain, & N. Radziwill, 11 March 2009, SPIE Newsroom. “Web service automates access to multiple astronomical data sources.” [http://spie.org/x34139.xml? ArticleID=x34139](http://spie.org/x34139.xml?ArticleID=x34139)


Usage Logs

A number of NVO participating organizations have implemented standard interfaces to their web and service logs, and we collect these logs in order to track VO-enabled use of data and services. The main NVO web site continues to be used frequently, with a new all-time peak in October 2008 of nearly 662,000 hits and over 548,000 hits in November 2008. These peaks are probably a result of use and take-up of VO tools during and after the NVO Summer School. Web site usage declined in December and January to an unusually low level, but rose again in February and March (data for April is incomplete).


Monthly Usage of NVO Web Site (<http://www.us-vo.org>):




Monthly Hits at NVO Organizations Originating from NVO Applications:



Monthly Users of Any NVO Application or Website:



Cumulative Use of NVO Services:



Acronyms

AAS	American Astronomical Society
ADC	Astronomical Data Center
ADEC	Astrophysics Data Centers Executive Committee (NASA)
ADQL	Astronomical Data Query Language
AIPS++	Astronomical Image Processing System++ (NRAO)
API	Applications Programming Interface
AVO	Astrophysical Virtual Observatory
CACR	Center for Advanced Computational Research (Caltech)
CADC	Canadian Astronomy Data Centre
CDS	Centre de Données astronomiques de Strasbourg
CMU	Carnegie Mellon University
CXC	Chandra X-Ray Center
CY	calendar year
DAG	Directed Acyclic Graph
DAGMan	Directed Acyclic Graph Manager (Condor)
DAML	DARPA Agent Markup Language
DARPA	Defense Advanced Research Projects Agency
DIS	Data Inventory Service
DM	Data Model
DOE	Department of Energy
DPOSS	Digitized Palomar Observatory Sky Survey
DTD	Document Type Description
EPO	Education and Public Outreach
ESTO	Earth Science Technology Office (NASA)
ESTO-CT	ESTO Computational Technologies (NASA)
FIRST	Faint Images of the Radio Sky at Twenty Centimeters
FITS	Flexible Image Transport System
FNAL	Fermi National Accelerator Laboratory
FTP	File Transport Protocol
FY	fiscal year
GB	gigabyte
GLU	Générateur de Liens Uniformes (uniform link generator)
GRB	Gamma Ray Burst
GriPhyN	Grid Physics Network
HEASARC	High Energy Astrophysics Science Archive Center
HTTP	HyperText Transport Protocol
IPAC	Infrared Processing and Analysis Center (Caltech)
IRAF	Image Reduction and Analysis Facility (NOAO)
IRSA	Infrared Science Archive (IPAC)
ISI	Information Sciences Institute (USC)
ITWG	Information Technology Working Group (NASA data centers)
iVDGL	International Virtual Data Grid Laboratory
IVOA	International Virtual Observatory Alliance
IVORN	International Virtual Observatory Resource Name

JDBC	Java Data Base Connectivity (Sun, Inc., trademark)
JHU	The Johns Hopkins University
MAST	Multimission Archive at Space Telescope (STScI)
MB	megabyte
MOU	Memorandum of Understanding
MWG	Metadata Working Group
NASA	National Aeronautics and Space Administration
NCSA	National Center for Supercomputer Applications
NED	NASA/IPAC Extragalactic Database
NESSSI	NVO Extensible Secure Scalable Service Infrastructure
NOAO	National Optical Astronomy Observatories
NPACI	National Partnership for Advanced Computational Infrastructure
NRAO	National Radio Astronomy Observatory
NSF	National Science Foundation
NVO	National Virtual Observatory
OAI	Open Archives Initiative
OASIS	On-line Archive Science Information Services (IRSA)
OGSA	Open Grid Services Architecture
OIL	Ontology Inference Layer
OWL	Web Ontology Language
PB	petabyte
PMH	Protocol for Metadata Harvesting (of OAI)
Q	quarter
QSO	Quasi-Stellar Object
RC	Replica Catalog
RDF	Resource Description Framework
REST	Representational State Transfer
RLS	Replica Location Service
ROME	Request Object Management Environment
SAO	Smithsonian Astrophysical Observatory
SAWG	Science Archives Working Group (NASA)
SAWG	System Architecture Working Group (this project)
SciDAC	Scientific Discovery through Advanced Computing (DOE)
SDSC	San Diego Supercomputer Center
SDSS	Sloan Digital Sky Survey
SDT	Science Definition Team
SIAP	Simple Image Access Protocol
SOAP	Simple Object Access Protocol
SRB	Storage Resource Broker
SSAP	Simple Spectral Access Protocol
STScI	Space Telescope Science Institute
SWG	Science Working Group
TAP	Table Access Protocol
TB	terabyte
UCD	Unified Content Descriptor
USC	University of Southern California

UDDI	Universal Description, Discovery, and Integration
UIUC	University of Illinois Champaign-Urbana
USNO	United States Naval Observatory
USRA	Universities Space Research Association
UWS	Universal Worker Service
VDL	Virtual Data Language
VDS	Virtual Data System
VO	Virtual Observatory
VO	Virtual Organization
VOQL	Virtual Observatory Query Language
WBS	Work Breakdown Structure
WebDAV	Web-based Distributed Authoring and Versioning
WSDL	Web Services Description Language
XML	Extensible Mark-up Language
2MASS	Two-Micron All Sky Survey

